

In the Matter of)
)
)
 Protecting and Promoting the Open Internet) GN Docket No. 14-28
)
)

We therefore suggest that the Commission adopt a framework based upon the one adopted by the Computer Inquiries, using a layered approach. This would apply common carrier regulation and open access to the lower layers of access networks, allowing unfettered competition and market discipline to return to the retail ISP marketplace, while the Internet itself—like all information services—would remain

essentially free of Commission regulation. This would allow ISPs the freedom and flexibility to evolve and market their services, while allowing customers the corresponding freedom and flexibility to move to a different provider of Internet service over the same physical medium.

Contents

Introduction and Executive Summary.....	1
Restoring a layered approach to regulation.....	3
Horizontal components	4
Vertical layers	5
Competitive forces are weakest in the access networks.....	9
Identifying what to regulate	10
Incumbents did not behave as forecast	10
Best-effort service is not always ideal for innovation.....	11
Congestion control is an ongoing problem	12
Fast and slow lanes are the wrong metaphor	13
Polyservice networks largely lack interconnectivity.....	14
Non-Best Effort networks introduce blocking	15
A simple, specific rule would be the optimal solution.....	16
Specific questions raised in the Notice	19
Section 706 has no teeth	19
Edge providers should have no special privileges	21
Edge providers are large-scale customers of ISPs	21
Edge providers should not be singled out for special treatment	23
Specialized services are beneficial and fundamentally harmless.....	24
ISPs should not need to go to great lengths to create meaningless metrics	25
Blocking and differentiation are both a vital part of Internet management	26
Priority is an oversimplification.....	28
Competition is woefully inadequate	29

A simple solution leads to maximum openness	30
Collusion is a structural issue that needs addressing	33
Section 706 should be construed narrowly	34
Title II should be applied, but narrowly.....	35
Reclassification is fully justified by the factual record	38
The open Internet is today’s free press	40

Restoring a layered approach to regulation

There are essentially two different ways to analyze a network. One is to look at its components horizontally, like beads on a string. Each component—a customer, a router, a server, etc.—is thus one bead, and the string is the telecommunications facility that links it to other beads. The other approach is vertical, to define layers that represent a functional decomposition of the activities taking place at each point in the network.

The historic public telephone network offered a single service across a single wire. It was thus easily analyzed without layers. But in the Computer Inquiries, beginning in the 1960s, the Commission recognized the importance of layering. We are not referring, in this case, to the textbook OSI Reference Model, which is primarily of academic interest, or the Internet’s corresponding protocol model, but to a more functional model that reflects practical, legal, and regulatory reality. The Computer Inquiries used layering to create a bright line between the lower-layer functions that were subject to regulation as basic services and the higher-layer functions that were deemed enhanced services. Computer Inquiry rules then imposed specific restrictions upon the behavior of local exchange carriers in order to ensure that their monopoly power over basic services did not interfere with competition among third party providers of enhanced services.

A solution to the Open Internet problem requires both horizontal and vertical analysis. Neither approach is sufficient on its own. Horizontal analysis alone risks harming the Internet itself by regulating activities that do not need regulation, and which are statutorily information services and thus outside of the scope of Title II. Vertical analysis alone ignores the vast differences in market power between service providers at different points in the Internet ecosystem.

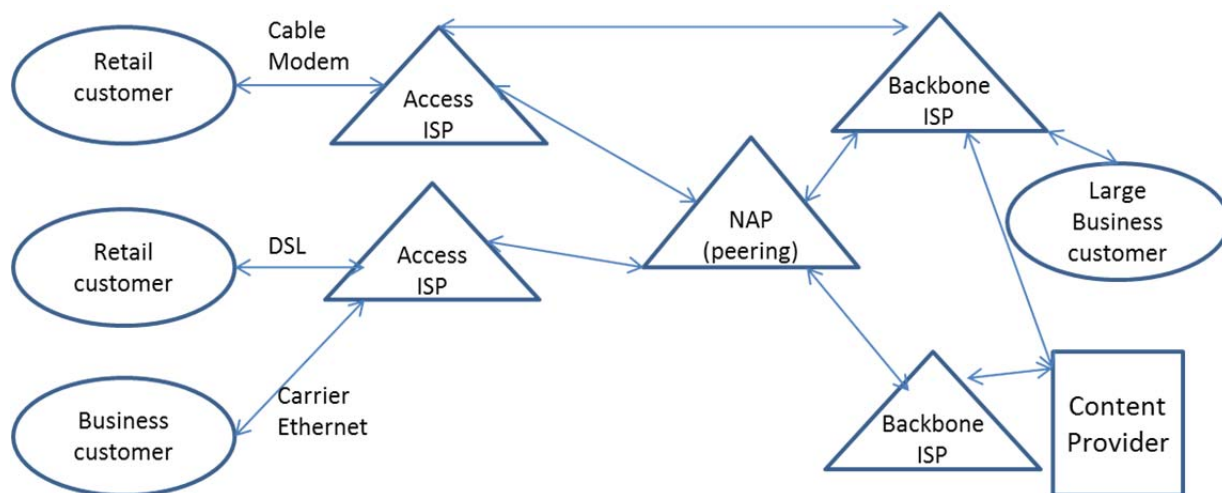


Fig. 1. Horizontal “beads on a string” model of Internet access.

Horizontal components

Fig. 1 is a simplified illustration of the horizontal components of a typical Internet connection within the United States. The retail customers are using what is typically referred to as “broadband” services, such as DSL, cable modems, and fiber to the home. The ISP who serves them directly can be referred to as an access ISP. Some business customers also connect to the same access ISPs, but in addition to the low-priced service options sold to residential users, business subscribers also use higher-priced access media such as Carrier Ethernet and Special Access circuits. Of these access facilities, only Special Access is currently subject to Commission regulation.

Access ISPs use backbone ISPs to reach most sites outside of their own networks. In practice, the major access ISPs nowadays are affiliated with backbone ISPs, but the roles are distinct. Backbone ISPs generally peer with each other at Network Access Points (NAPs), as well as directly (not illustrated), and access ISPs who have facilities to NAPs can peer with, or purchase upstream service from, multiple backbone ISPs. The circuits between these ISP locations are conceptually similar to the circuits used by long-distance telephone providers (IXCs) between their Points of Presence. They are competitive between major cities, though the middle mile connections from the backbone to access ISPs in some rural areas can be very expensive.

Content Distribution Networks (CDNs) are computer services that provide caches of popular content at or near access ISP core locations. This enables access ISPs to retrieve content without utilizing backbone facilities resulting in both greater bandwidth efficiency and improved service delivery.

Content providers are generally customers or, in the case of the largest few, peers of backbone ISPs. They inject more packets than they receive, the opposite of the retail broadband subscribers who are their users. Thus there is a net flow of packets from the content provider to the backbone ISP to the access ISP to the subscriber. If the content is cached on the CDN, the flow is shortened, though the CDN cache is frequently refreshed across the backbone. In either case, content providers have good competitive options (except to the extent that a given subscriber's service provider is the bottleneck). They generally collocate their servers in carrier hotels or data centers that are directly served by multiple backbone network providers in order to essentially bypass the access ISP role for their service delivery.

A key issue for the Commission to deal with is the relative competitiveness of each of these links. If a link is available on a highly-competitive basis, lighter regulation is warranted. Access circuits are generally only available on a monopoly or duopoly basis—most consumers can only get broadband Internet service from the incumbent telephone company or cable company. Other access media exist, but are generally not competitive. CMRS access provides mobility, and is geared towards smartphones; with typical data rates in the \$10-15/GB range, it is orders of magnitude too expensive to compete with DSL or cable for what is becoming typical consumer usage in the 30-100 GB/month range. Satellite, also usually capped due to capacity constraints, suffers from high latency, which makes it virtually unusable for many applications including VoIP and gaming. Wireless ISPs provide good service in many rural areas but the spectrum cannot support WISP access to a very large share of urban customers. Some WISPs even refuse to provide service to customers who can get cable, because they know that cable offers a more cost-effective solution and thus subscribers who choose the WISP are likely to have had cable service shut off for non-payment. Thus taken as a vertically-integrated whole, the access link is the least competitive part of the horizontal model.

Vertical layers

A layered analysis begins by seeking to define the relevant layers. This should be done in a manner that is both easy to administer and technologically neutral. Another constraint is the statutory definitions of *telecommunications service* and *information service*. There is nothing wrong with these categories *per se* but it must be noted that while their definitions reflect an understanding of the state of the art in the mid-1990s, their specific language also reflects the state of the art in the late 1970s when *Computer II* was being deliberated. Thus the language of the law needs to be viewed in terms of its clear intent. The Computer II concepts of *basic* and *enhanced* services still ring true, even if most (if not all) of the enhanced services of the day have merged into today's Internet.

We suggest that the clear intent of the distinctions in both Computer II and the Telecom Act definitions is the distinction between carriage and content, wherein *internetworking*, the voluntary exchange of information traffic between network operators for their mutual benefit, was clearly meant to be on the content, or information, side of the line. The intent was to regulate the basic services (telecommunications) being provided on a monopoly basis, and to not regulate activities done *with* telecommunications, but rather to facilitate them by ensuring access *to* telecommunications.

Computer II achieved this via structural separation. Its key tool was to prohibit local exchange carriers from offering enhanced services unless they were provide on a structurally-separated basis. Computer III modified this to allow for accounting separation, while still requiring basic services used in its provision of enhanced services to be made available to other enhanced service providers.

Language in Computer II that specified certain activities that constituted enhanced services encompassed the preponderance of contemporary computer communications, while explicitly carving out voice telephony, which remained basic. The Commission had just adopted the *Exchange Network Facilities for Interstate Access* (ENFIA) rules in 1979, formalizing contributions from competitive long-distance carriers, and did not want the simple patching of telephone calls together to escape the new access charges. Other than protecting access charges (an obsolete concern going forward), it was fundamentally a straightforward division between carriage and content.

We thus suggest that for the purposes of this proceeding, the following layered model be used.

Content	<i>Customer content</i>
Applications	
TCP/IP (or alternative) transport	<i>Information service</i>
Lit telecommunications (DSL, DOCSIS, PON, etc.)	<i>Telecommunications [service]</i>
Wire, cable, and fiber; CMRS facilities	
Conduits, ducts, and poles	<i>Network elements</i>

Figure 2. Functional layers and their constituents.

At the foundation of this stack are the network elements, facilities that by themselves provide no service but create the physical transmission infrastructure that is essential to the provision of services. Provision of these on an unbundled basis is the subject of Section 251 of the Telecommunications Act and numerous Commission proceedings. It is clear that whether or not an element should or shouldn't be

unbundled has little to do with the payload it carries and much to do with whether or not functional substitutes are available. The Commission has not reviewed this in depth since the 2003 *Triennial Review Order* and the 2004 *Triennial Review Remand Order* and it is generally outside of the scope of this proceeding.

Telecommunications service includes any number of technologies that “light” the network elements, multiplex them, and make them useful at specific bandwidths or bit rates. So-called broadband services that fall into this layer include “raw” DSL¹, the DOCSIS service provided using cable modems, and the Passive Optical Network layer used by some fiber services. These services should in all cases be provided on a truly bit-neutral basis, and their payloads should not be restricted to or defined by any specific set of applications or protocols. This need not, however, constrain content-neutral traffic-management options, like those historically applied to Frame Relay and Asynchronous Transfer Mode (ATM) services, which would fit into this layer.

Information services are the payload of telecommunications services. In the specific context of the Internet, the entire TCP/IP protocol stack, beginning at what is referred to as the Internet Protocol (IP), is an information service.² Other software-based protocol suites, while not currently in such widespread use, can also constitute information services; there is no magic pixie dust in IP itself. MPLS, a shim layer typically used just below IP, has been provided as both a telecommunications service and an information service; there is no contradiction in offering it both ways. We also note the development of RINA (Recursive InterNetworking Architecture), which is currently being prototyped in Europe by the IRATI (Investigating RINA as an Alternative to TCP/IP) project³, which can potentially be applied in the delivery of telecommunications, information, and cable services⁴. The ability to innovate in this manner is dependent upon access to telecommunications, not to IP itself.

It is not easy to separate out the application from the underlying layers of the information service. In the context of the Internet, the TCP/IP protocol suite itself is not truly layered. That is, the operation of one layer is not independent of all non-adjacent layers, a principle understood by 1980 when the OSI program

¹ This is still provided under tariff by some rate-of-return ILECs, usually under the terms of NECA Tariff 5, Section 8.

² This does not mean that all uses of IP are information services. IP is fundamentally just a multiplexing header; it can be used to provision telecommunications services as well. This applies, for instance, to some fixed IP-based telephone services provided by facilities-based operators, such as LECs and cable providers. It can also be used in the provision of a cable service, as in a cable system using IPTV instead of RF channelization.

³ <http://irati.eu>

⁴ RINA recurses the same few protocols in as many layers as needed, rather than using separate protocols at different layers. Each instance of the protocols may vary in scope and policies.

was under way. TCP/IP dates back to 1974⁵, and while it has had many patches since then, its fundamental architecture still reflects certain expediencies taken in that time frame by the ARPANET development community. For example, the IP address theoretically belongs to the IP layer, but is sometimes used by the application layer; names are resolved by DNS inside the application, not the network. Since an IP address simply identifies a single point of attachment to a given computer, it is clearly inadequate to represent any kind of popular “edge provider” service, which is typically hosted on a large number of computers, often in many locations. Thus the application may invoke a process of redirection in order to reach a specific server that has the resources to meet its request. This may involve communication with underlying network resources. Network address translation is also a common ISP function that requires the network to modify some application layer protocol headers⁶. For this reason, we view the application as being at least partially within the underlying information service, and partially within the customer’s content. And the line between these two layers is not a bright one, but, in practice, a somewhat porous boundary that is also an area of innovation and creativity.

The ISP customer’s content—application processes and data located on host computers—is at the top of the stack. It is clearly not part of telecommunications, but it is the resource in which most users of the Internet are most interested.

We note that within the context of the TCP/IP stack, the layer *below* IP, which is external to the stack itself, is often referred to as the *local network*, or simply the *network*. IP is thus an *internetwork* protocol because it carries packets across multiple, heterogeneous, separately-controlled networks.⁷ Using those definitions, the Commission can enforce *network* neutrality (on telecommunications) without imposing new regulatory burdens on the *internetwork* (the Internet).

Having identified the relevant layers, we now have part of the framework needed to identify the potential market failures that may call for regulatory response, and, just as important, to identify those that need not be subjected to regulatory interference.

⁵ TCP/IP’s public debut was in the article *A Protocol for Packet Network Intercommunication* by Vinton G. Cerf and Robert E. Kahn, IEEE Transactions on Communications, Vol. COM-22 No. 5, May 1974. IP was deemed a separate protocol as of TCP version 4, in 1978, hence IP version 4.

⁶ This would not be the case if applications did not put IP addresses in their own headers, a dangerous but nonetheless common practice. But other than that, network address translation (NAT) is a harmless process which merely swaps one connection ID (IP address + port number) for another. It is a beneficial process when used for security purposes, and is thus the norm in both home and corporate firewalls. Some ISPs use it in order to conserve IPv4 addresses.

⁷ Postel, J., IEN54, *Internet Protocol Specification, Version 4*, September 1978.

Competitive forces are weakest in the access networks

The strong public demand for “network neutrality” arose in direct response to the Commission’s 2005 *Wireline Broadband Order*⁸. Since then, DSL has been treated as a vertically-integrated service, not layered. Its content is access to the Internet. The Internet is an information service, and the entire facility that carries it is become treated as an information service. As Steven Hawking retells in *A Brief History of Time*, a “little old lady” responds to an astronomer’s lecture with disdain:

“What you have told us is rubbish. The world is really a flat plate supported on the back of a giant tortoise.” The scientist gave a superior smile before replying, “What is the tortoise standing on?” “You’re very clever, young man, very clever,” said the old lady. “But it’s turtles all the way down!””⁹

Since information services are not subject to Title II common carriage requirements, that reclassification excluded third-party ISPs from the right of access to these facilities. Some ISPs who had existing arrangements were allowed to continue them for a time under commercial agreement, but even that was temporary, as we will discuss in more detail below.

The obligation of LECs to provide common carriage to competing ISPs was a direct consequence of the Computer Inquiries. The cable industry had never been subjected to the same requirements, and in its 2002 *Internet over Cable Declaratory Ruling*¹⁰, also known as the *Cable Broadband Order*, the Commission concluded that cable modem service, “as it is currently offered”, is properly classified as an interstate information service. It subsequently did not choose to require cable operators to offer access to third parties; few did so voluntarily. Likewise the Commission had exempted FTTH from Computer II/III requirements and most Section 251 unbundling. However, the availability of DSL to a majority of the public prior to the Wireline Broadband Order did create the opportunity for third-party non-facilities-based ISPs to address the mass market, and thus the choice of ISPs for consumers was greater than the number of facilities owners (i.e., one or two within each local context). This dynamic, coupled with the open entry to the ISP business made possible by the Computer Inquiry rules, created and sustained the environment of an open Internet, if only by the threat that if the facilities owners did something to sufficiently displease their customers, an alternative could arise.

Today’s cable/ILEC duopoly is assumed to provide a choice of two broadband service providers to most consumers. But there are many areas in which only one is available, and a mere handful in which an

⁸ *Appropriate Framework for Broadband Access to the Internet over Wireline Facilities*, Report and Order and Notice of Proposed Rulemaking, FCC 05-150, in CC Docket 02-33.

⁹ Hawking, Stephen, *A Brief History of Time*, 1988, Bantam Books; p. 1.

¹⁰ FCC 02-07, at 7.

overbuilder, such as Google Fiber, RCN, or Wide Open West, provides a third option. Overbuilding, however, is now very rare, with operators such as RCN operating bankruptcy assets originally financed by massive losses on the part of earlier owners.

Intercity assets, on the other hand, are highly competitive. Long-haul fiber is owned or controlled by a number of backbone service providers, including Level 3, Verizon, AT&T, CenturyLink, and XO. Entry into the backbone ISP market is not closed—transmission facilities are available competitively from those firms and other IXC's, while Allied Fiber is building new dark fiber facilities for wholesale customers. If there is a market failure, or if there is a reasonable likelihood that market failure could occur, then it is going to happen in the access networks, not the long haul networks.

Identifying what to regulate

Given the horizontal and vertical analyses presented above, one can decompose the Internet system into a number of components. This table briefly characterizes them:

	Access	Regional/aggregation	Backbone
Customer content	end users	CDN, small edge providers	large edge providers
Information service	easily separable from lower layers	generally competitive	competitive
Telecommunications	monopoly/duopoly <i>recommended focus</i>	competitive in most markets	competitive
Network elements	monopoly/duopoly (limited; for CLECs only)	limited availability	limited availability

The statutory basis of the Commission's authority to regulate is clearly strongest in the lower layers, especially Telecommunications services. From a market failure avoidance perspective, its focus should therefore be on the intersection of telecommunications with access. At this level the backbone and the regional/aggregation entities, such as CDNs and data centers, are generally competitive; and if independent ISPs were granted the right to use existing local access facilities to reach their subscribers from their regional aggregation points, there would likely be a competitive market for ISP services, as there was during the dial-up era.

Incumbents did not behave as forecast

In the *Wireline Broadband Order*, the Commission justified removing Computer II/III requirements from LECs in part by stating its assumption that they would continue to provide wholesale service to ISPs even without the obligation to do so:

At this time, facilities-based wireline carriers are the only providers of broadband Internet access services that are compelled by regulation to make such an offering available. As stated above, this compulsion is not the result of the Commission's analysis of broadband Internet access services specifically, but rather is the product of the application of legacy rules adopted decades ago. Therefore, we cannot state unequivocally that incumbent LECs would not otherwise provide wholesale access, absent this compulsion. In fact, the record shows that incumbent LECs would and indeed already do provide such access, albeit through arrangements other than a mandatory tariff regime that requires a standardized general offering.¹¹

The same Order predicted that cable operators would voluntarily do the same, although in practice such arrangements had grown quite rare by 2005:

For their part, cable operators, which have never been required to make Internet access transmission available to third parties on a wholesale basis, have business incentives similar to those of incumbent LECs to make such transmission available to ISPs, and are continuing to do so pursuant to private carriage arrangements.¹²

Had these forecasts been accurate (and it is not even clear that they were sincere), then there would be no issue, as there would still be a wide choice of mass-market ISPs. But neither the carriers nor the cable operators made mass-market wholesale access available on a widespread basis, and its vestiges are in decline. This dramatic difference between forecast and reality alone should provide the Commission with the necessary factual predicate to revisit its positions.

Best-effort service is not always ideal for innovation

The fundamental service offered by the public Internet today is often referred to as “best effort.” This characterization is not a pejorative reference to a network that “doesn't try hard enough”; it refers to the fact that the Internet core is deliberately designed to be a transparent fabric of routers that forward packets along the best paths available from a source to a destination, without the additional overhead of error detection and recovery on each individual hop. Some packets make it to their destination; some are lost. The mechanisms necessary to recover lost packets—if necessary—are implemented in systems attached to the network, not inside the network itself.

¹¹ *Wireline Broadband Order*, FCC 05-150 at 63.

¹² *Id* at 64.

Congestion control is an ongoing problem

It is sometimes assumed that higher “bandwidth”—a higher bit-per-second capacity along links—will prevent packet loss. This is simply not true, and it is easy to illustrate. The typical personal computer nowadays has an Ethernet jack that operates at 100 or 1000 Mbps. The typical content-provider server has an Ethernet jack that operates at 1000 or 10,000 Mbps. The typical home subscription has an upstream rate cap of no more than 10 Mbps and a downstream cap in the 4-100 Mbps range. In either direction, full-speed streaming is not possible.

The signal that enables senders and receivers to modulate this flow of data—to perform the required speed matching—is packet loss, detected by the TCP process (or the application itself if TCP is not used) running on the sending computer. If a packet is not acknowledged on time, the sending computer assumes that it was discarded due to congestion. It then lowers its transmit window—the number of unacknowledged packets allowed to be outstanding in a given TCP flow—to one, and increments this number as packets are acknowledged. This process is known as slow start, and was introduced to TCP in 1987 after the ARPANET suffered congestion collapse. That phenomenon occurred when lost packets were being retransmitted, increasing the total load on the network, increasing packet loss, leading to a very busy network with very little useful data getting through. Slow start was added to TCP, which runs in the host computer systems, not the network itself; the cooperation of network users was-, and still is-, thus required. Slow start allowed speeds to scale up, so that TCP adapts automatically from slow modem links to high-speed LANs.

This worked well in the ARPANET and subsequent government-funded networks, where non-cooperation could at least theoretically result in one’s removal from the network. In the public Internet, there is no such mechanism. It has however been an ISP’s prerogative to modulate or even block flows that it deems harmful to the network. This is one example of where the design most certainly did not assume common carrier-like “neutrality” of all IP packets.

Another issue arises from streaming, where TCP may or may not even be utilized, but demand is relatively inelastic. Streaming over UDP—necessary for real-time applications such as VoIP—does not back off at all in the face of packet loss. Some streaming video systems however use TCP with slow start¹³ to incrementally upload a few seconds’ worth of video at a time. Lost packets can be retransmitted before the receiver’s buffer finishes displaying its contents, and the video quality can be adjusted

¹³ TCP is a protocol; slow start is a policy. It is possible to use TCP without slow start, or with a variant policy. It is even theoretically possible to use TCP without retransmission for streaming, though it is rarely done, as UDP is a simpler alternative.

downwards, if necessary, to continue gap-free operation. Notable here is that TCP video streaming can recover from packet loss, while VoIP does not, and thus packet loss degrades the latter's quality.

Congestion, then, is not a problem of insufficient capacity. It is a design feature of TCP/IP. Because some applications—especially real-time applications that do not tolerate loss well—are harmed by it more than other applications, a monoservice “best effort” network *by its very nature favors some applications over others*. Expanding the capabilities of the Internet requires looking beyond the monoservice model of a purely “neutral” network that treats all packets equally. There are however alternative approaches to defining nondiscriminatory behavior.

Fast and slow lanes are the wrong metaphor

The common description of a non- neutral internet¹⁴ is one that has both fast and slow lanes. This analogy to road traffic is not appropriate. We suggest that the distinction instead be made between a monoservice network, such as the present best-effort Internet, and a polyservice network. Polyservice networks do exist today. Many enterprises, for instance, make use of MPLS services that allow for customer-defined prioritization. The fastest-growing segment of the wireline transmission industry is probably Carrier Ethernet, a polyservice offering that allows each connection to have a defined set of Quality of Service (QoS) parameters. These typically include a Committed Information Rate (CIR) and Excess Information Rate (EIR), each rate averaged over defined burst sizes. CIR is the rate at which traffic is carried at top priority, assumed to be lossless in normal operation. EIR is the amount of discard-eligible traffic, above the CIR, which is carried on a discard-eligible (i.e., “best-effort”) basis. Above the EIR, traffic is discarded at the ingress.

In such a network, “fast” and “slow” are inappropriate metaphors. Take, for instance, a very practical access option for many locations. This may include a connection to a telephone service provider that has a CIR of, say, 1 Mbps, and no EIR, and a connection to an ISP with no CIR but an EIR of 50 Mbps. The ISP connection is faster but tolerates loss, since there is little point in being lossless when the Internet itself is lossy. But the slower telephone connection has a higher quality, and is thus capable of supporting impairment-sensitive voice-grade services such as fax, modems, and security alarms, as well as supporting more natural voice conversation.

None of this is new. Polyservice networks were being designed in the 1980s. The original specification for IP was written in 1978 and documented in IEN54. The header includes four precedence bits,

¹⁴ Upper-case “Internet” is the current global TCP/IP network. Lower-case “internet” is a generic term; there can in theory be multiple internets. The Internet is essentially a prototype internet. Private internets also exist, some with firewalled connections to the public Internet. Interisle manages one such network on behalf of a group of local public safety agencies.

described as “Specifies one of 16 levels of precedence.”¹⁵ In 1998, RFC2474¹⁶ replaced the admittedly little-used precedence field with the Differentiated Services field. The first sentence of that RFC’s text is quite clear about its purpose: “Differentiated services are intended to provide a framework and building blocks to enable deployment of scalable service discrimination in the Internet.” DiffServ has seen significant usage, though more on enterprise networks and within individual service providers’ networks than across provider boundaries (peering) on the public Internet. This discrimination is *not* insidious or intended to create second-class users. Rather, it is a tool for *managing network performance to improve the utility of the TCP/IP protocols*.

Even without explicit precedence, TCP/IP’s mechanisms create a significant, if inadvertent, degree of discrimination. The aforementioned slow-start mechanism ratchets up the window size based upon acknowledged packets, each increment taking essentially one end-to-end round-trip across the network. Thus a path that takes 200 milliseconds for an acknowledgement to reach the sender will increase its speed much more slowly than one that takes only 50 milliseconds. This is one advantage of content delivery networks: by being fewer hops and fewer miles away from the user, they can cycle through slow-start more rapidly, achieving a higher speed before congestion causes a packet drop. TCP thus favors short hops over long, in marked contrast to rate-based services.

Polyservice networks largely lack interconnectivity

Polyservice offerings today are largely confined to individual owners’ or providers’ networks. It is quite common to purchase an MPLS service from one ISP that supports QoS and traffic isolation; it is not yet common for such services to be offered on an internetwork basis. This is both a technical and a business issue. In technical terms, MPLS and Carrier Ethernet do not have strong network-to-network interface (NNI) definitions. Frame Relay and ATM, polyservice telecommunications service offerings of the 1990s, had both a user-network interface (UNI) and an NNI in their standards. Due to restrictions on interLATA services by Regional Bell Operating Companies, nationwide networks were composited by concatenating RBOC intraLATA networks to an IXC’s interLATA offering via an NNI. With the subsequent grant of Section 271 interLATA authority to all of the Bell companies, this distinction no longer had to be made, and the surviving national carriers prefer to offer end-to-end services. But this need not be the only option going forward, and it is certainly not desirable for these enterprise-critical services to be banned, or mass-market consumers to be deprived of their benefits, because they are not “neutral.”

¹⁵ Postel, J., Sept. 1978, IEN54, *Internet Protocol Specification, Version 4*, at 16.

¹⁶ Dec. 1998, RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

Given the increasing demand for audio and video services, and the need for edge providers to have connectivity across multiple access and transit networks, the time may be ripe for a renewed emphasis on polyservice interconnectivity. This will necessarily include both financial and technical arrangements. The financial arrangements of polyservice peering will need to be worked out in the marketplace. Attempts to dictate them prospectively, by treating them as regulated interconnection along the lines of switched access service, are likely to fail. Only the presence of monopoly power at choke points in the network, such as mass-market access circuits, may require special scrutiny. And that can be addressed as well by applying a layered approach.

Non-Best Effort networks introduce blocking

While the term “blocking” is used in the Notice to refer to intentionally preventing connectivity to an end point, the same word is also a term of art used in telecommunications networks, such as the PSTN. It refers to the inability of the network to service a request due to momentarily-inadequate network capacity. Circuit-switched networks are traffic-engineered to stay under a maximum blocking percentage during the busy hour. The standard formulas for PSTN traffic engineering make the assumption that traffic demand is Poissonian—the likelihood of an additional call attempt is independent of the number of calls already in progress—but also predictable based on historic patterns. Telephone call demand changes relatively slowly, by Internet standards; and each telephone call uses exactly the same resources (64000 bits per second, and associated control-plane resources) as every other.

The multimedia streaming services that make us of the Internet today differ noticeably from the PSTN. For one, the demand generated by each connection varies; video may use multiple megabits per second, compared to about 90,000 bits per second for high-quality VoIP. For another, the network has no traffic engineering. Streaming takes advantage of the fact that many Internet links are heavily overengineered, often running at less than 50% utilization even during peak usage periods. There may be congestion, but edge providers attempt to work around this by working with content distribution networks, which cache copies of content at many places in order to minimize the number of hops and ISPs involved in a given end user’s connection. Still, there are no promises made, and there cannot be—a best-efforts network does not make promises of packet delivery.

TCP applications are extremely elastic. As noted above, the slow-start algorithm modulates the flow of TCP applications so that they automatically adjust to available capacity. This is the opposite of the PSTN, where demand is fixed. TCP applications automatically adjust to the bottleneck in the connection, whether it be peering or access, and whether it be broadband, mobile, or even dial-up.

If a network offers service levels other than best efforts, for example a connection with a committed information rate, then it necessarily is subject to blocking. If a given link has 100 Mbps of committable capacity and it is all committed, a request for another 5 Mbps of committed information rate should not be filled. The request is thus blocked (unless an alternative path is available). The connection-request protocol must thus return some kind of failure message (equivalent to fast busy on an analog PSTN line).

We point this out not to alarm, but to note that such blocking is almost inevitable, and probably better than the alternative. Absent a polyservice network, service to *all* users degrades when traffic demand exceeds supply. If TCP backs off and streaming doesn't, TCP applications suffer disproportionately. Real-time streaming (e.g., VoIP) suffers, though delayed streaming (e.g., one-way video) can adapt and seize a greater percentage of the total available capacity. This is technically "neutral" but not as functionally "open" as a polyservice network could be.

A simple, specific rule would be the optimal solution

Given the above, we suggest a simple and specific rule to preserve the open Internet: that the Internet itself—the higher-layer functions—not be generally regulated, but that regulation focus instead on open access to the lower-layer access network, where it is most needed and likely to be effective. The Commission should recognize that the forbearance granted from the Computer II rules, and its decision to exclude cable from similar regulation, is no longer appropriate. The mere fact that the network neutrality kerfuffle erupted within months of the 2005 *Wireline Broadband Order*, and not just coincidentally, is strong evidence that that Order was overly optimistic, and should be revisited.

This was clearly suggested by the D.C. Circuit Court of Appeals in its *Verizon* decision. That decision recounts the history of the Computer II ruling and how the definitions of telecommunications service and information service in the Telecommunications Act are based upon Computer II's own definitions of basic and enhanced services. The Court cited Justice Breyer's concurrence in *Brand X* that the FCC's position with regard to cable modems was at the very edge of Chevron deference ("concluding that the Commission's decision to exempt cable broadband providers from Title II regulation was 'perhaps just barely' within the scope of the agency's 'statutorily delegated authority'")¹⁷. The Court is clearly, then, amenable to a Title II solution based upon a layered approach, applicable to both LEC and cable.

This regulation must be carefully focused. Horizontally, it should focus on the access links, on the left side of Fig. 1 above, as these are the least competitive and are the clear focus of concern about the openness of the Internet itself. Vertically, it should begin at the bottom of the protocol stack and extend no higher than necessary. In general this would be the telecommunications layer.

¹⁷ *Verizon* at 24.

The proposed rule, then, is simple:

Service providers should be required to make the telecommunications layer of their networks available to any requesting party on a common carrier basis, subject to Title II regulation, especially Sections 201, 202, 208, and 254.

This telecommunications service should be *bit-neutral* carriage which can be used to provide Internet access, but which can also be applied to other applications and higher-layer protocols, inasmuch as the payload of common carriage belongs to the customer, not the carrier. Examples of telecommunications protocols that can be used to fulfill this obligation without any involvement with the IP function whatsoever include DSL, DOCSIS, Carrier Ethernet, GPON, and LTE.

If the service provider is unable to make this lower layer service available, or chooses not to do so, then its common carrier obligation extends up the protocol stack to the lowest protocol layer that it makes available as a service boundary.

This might frequently be IP. In such cases, its interconnection¹⁸ at the regional point of aggregation should also be subject to common carriage. This IP service should be configured in a manner that permits other ISPs to “tunnel through” it.¹⁹ Routine content-aware IP network management procedures, including spam filtering, would not be performed on these wholesale links; they would be left to the service provider’s customer (typically an ISP). Traffic shaping, however, could be performed pursuant to the terms of a tariff or generally-available contract, as a bit-neutral managed service.

Not all Internet service providers should be subject to the proposed rule. Not all ISPs have equal market power. The majority of mass-market broadband services today are offered on a vertically-integrated basis by cable companies and incumbent local exchange carriers. Mobile Internet access service is also provided by licensed Commercial Mobile Radio Service providers. These companies all own *rivalrous* facilities: Their use of physical facilities, or exclusively-licensed radio spectrum, excludes others from simultaneous use of the same facilities. In the case of telephone and cable companies, pole and duct space are rivalrous—the next carrier on a pole must typically pay for make-ready. Hence these service providers should be subject to the aforementioned Title II regulation of the lower layers of their access networks.

In contrast, ISPs that only lease facilities from carriers, or wireless ISPs who use unlicensed (Part 15) or non-exclusively licensed (e.g., Part 90 subpart Z) frequencies, do not make use of rivalrous facilities.

¹⁸ Interconnection is generally the term applicable to Title II services; peering is applied to the normally-unregulated activity performed by information service providers.

¹⁹ This would include transparency requirements plus a minimum packet size large enough to support the nested headers.

These non-carrier²⁰ ISPs should not be subject to open access obligations. ISPs affiliated with CLECs which are not affiliated with ILECs in a given location also should not be subject to these regulations, as they also lack market power.²¹

The regional point of interconnection (POI) for open access customers should generally be a major carrier hotel or other carrier-neutral location. Existing telephone company tandem offices are traditionally useful for Title II interconnection, but if service is offered only in IP form, then the point of interconnection should be something that is more attractive to ISPs. If the provider's service area and facilities do not extend to such a location (as might be the case with a rural carrier or cable company), it should not be obligated to provide free backhaul to a distant city; it should make an aggregation point available in or near its service area, such as at a central office or head end, where backhaul facilities are available.

The purpose of this rule is to ensure that open ISP access to mass-market subscribers is restored, and applied equally to both cable and telephone companies. While the 2002 *Cable Broadband Order* upheld in the *Brand X* decision may have been correct at the time, market conditions have changed significantly since then. Telephone and cable companies are no longer the arch-rivals they once were; they have begun to collude and co-market each other's services. Telephone companies are abandoning their wireline plant in some areas, especially where states have relieved them of their carrier of last resort obligations. DSL has been grandfathered in some places. In other words, cable is in the process of becoming the new wireline monopoly in many places. Most cable companies have not yet crossed the threshold of Section 252(h)(2), which allows for ILEC designation to be transferred, but in many areas they are much closer to being ILECs than they once were. And certainly the rosy predictions of the 2005 *Wireline Broadband Order* have turned out to be wildly over-optimistic. The protections of the Computer II/III era are more important now than ever.

Our simple proposed rule reorients telecommunications regulation in a manner that supports the open Internet without directly imposing regulation upon the Internet itself. It restores the balance that allowed apparent network neutrality (customer-focused network management) to flourish as a natural phenomenon of the competitive market, not a perennial conflict involving lawyers and lobbyists. They also put the United States in line with other Western democracies, virtually all of which impose similar

²⁰ We note the terminology used in Canadian regulations, in which facilities-owning ISPs, generally cable and ILEC, are deemed *primary*, and their wholesale customers are *secondary*. This approximates the model we have in mind, clarifying that WISPs should also be viewed as equivalent to secondary providers and not subject to the obligations of primary service providers. This model avoids the philosophical question of whether or not a carrier is offering a telecommunications service or information service; it focuses instead on tangible physical facilities.

²¹ That is, the ILEC should not escape its obligations by transferring its DSL or FTTP activities to a CLEC affiliate, but it does not have the obligations where it is only a CLEC.

obligations on their telecommunications carriers in order to provide a competitive, lower-price, higher-performance Internet market.

Specific questions raised in the Notice

We now address a subset of the myriad of questions in the Notice of Proposed Rulemaking. These are in general addressed in the context of the above discussion, and thus we see no point in addressing specific issues that relate to matters that we do not believe should be subject to regulation. These include issues such as peering between backbone ISPs, access to the backbone by edge providers, and specific policies for managing traffic and applications within the Internet. Those issues can be managed by the marketplace, provided that no service provider gains monopoly power over any vital bottleneck facilities, and provided that routine antitrust enforcement remains a backstop against abuses. We note that the Commission's jurisdiction is plenary with regard to common carriage. With regard to content and higher-layer services, which are not common carriage, the Federal Trade Commission may also have authority, should abuses be noted.

Section 706 has no teeth

The Commission based its 2010 Order on Section 706 authority. The Court in its *Verizon* decision reduced the scope of those rules, while leaving the Section 706 justification intact for certain limited authority. But even this was derided as a “triple-cushion shot” and the Court seemed near its limit of *Chevron* deference. Nonetheless the Commission asks if new rules can or should be based on Section 706 rather than on its clearer Title II authority.

It appears that the Court and the Commission have apparently committed an error of law in accepting even the narrow form of the Commission's claim of Section 706 authority. It was not brought up by Verizon during its appeal, and thus was not shown to the Court, which could not have been expected to be aware of it; but it nonetheless risks making any further dependence on Section 706 futile. The *Verizon* Court noted that:

The Senate Report describes section 706 as a “necessary fail-safe” “intended to ensure that one of the primary objectives of the [Act]—to accelerate deployment of advanced telecommunications capability—is achieved.” S. Rep. No.104-23 at 50–51. As the Commission observed in the Open Internet Order, it would be “odd . . . to characterize Section 706(a) as a ‘fail-safe’ that ‘ensures’

the Commission’s ability to promote advanced services if it conferred no actual authority.” 25 F.C.C.R. at 17970 ¶ 120.²²

The Court, then, accepted the Commission’s use of the Senate report as its justification, interpreting legislative history to read certain power into Section 706. But this suffers from two fatal flaws. One is that the Senate Report that the Commission cited referred to Senate draft text that was *not* the final text of the law adopted by Congress. Section 706 was rewritten in conference. The Senate Report referred to a draft in which the last sentence of Section 706(b) read, “If the Commission’s determination is negative, it shall take immediate action under this section, and it may preempt State commissions that fail to act to ensure such availability.” But the actual law that Congress passed instead reads, “If the Commission’s determination is negative, it shall take immediate action to accelerate deployment of such capability by removing barriers to infrastructure investment and by promoting competition in the telecommunications marketplace.”

The reworded sentence removed a requirement that the Commission take direct action, and removed specific reference to preemption of state action. Instead it cited promoting competition, essentially the purpose of Sections 251-261. Even more critically, it removed “under this section.” That was the enabling clause, the legislative device that allows independent action under that law. Many other Sections of the Telecommunications Act contain that enabling clause; its removal from Section 706(b) essentially gutted it, turning it into mere general guidance.

A second flaw is that, unlike most of the Telecommunications Act of 1996, Section 706 is an outside section, not an amendment to the Communications Act of 1934. (Thus it was codified as 47 U.S.C. 1302, away from the Communications Act itself.) This was apparently not clear to the *Verizon* Court:

The Commission has identified at least two limiting principles inherent in section 706... First, the section must be read in conjunction with *other* provisions of the Communications Act, including, most importantly, those limiting the Commission’s subject matter jurisdiction to “interstate and foreign communication by wire and radio.” 47 U.S.C. § 152(a). Any regulatory action authorized by section 706(a) would thus have to fall within the Commission’s subject matter jurisdiction over such communications—a limitation whose importance this court has recognized in delineating the reach of the Commission’s ancillary jurisdiction. [emphasis added]²³

While Section 706 must be read in conjunction with the Communications Act, it is not itself part of that Act, which may further limit its authority, especially given that it lacks an enabling clause of its own.

²² *Verizon v. Federal Communications Commission*, No. 11-1355, at 24-25.

²³ *id* at 26.

We therefore suggest that the Commission set aside Section 706 and instead make judicious use of its clear Title II authority, not to regulate the Internet, but to regulate the fundamental telecommunications infrastructure upon which the Internet depends.

We do not believe that the Commission, as it noted, should build upon “the foundation begun under Chairman Powell, continued under Chairmen Martin and Genachowski, and reinforced by a decade of Commission policy.” Rather, it should build upon earlier policies that actually worked, that were less controversial, that were easier to administer, and that became the model successfully adopted by the rest of the world—only to be abandoned by the United States.

Edge providers should have no special privileges

The Commission raises questions about how edge providers, the operators of computer servers to which mass-market consumers frequently connect, should be seen. At 38, it asks:

We seek comment on the potential for, and development of, new business arrangements in the market between broadband providers and edge providers. What does the multi-sided market look like, and what are its effects on Internet openness? Do some types of broadband and edge provider arrangements (or aspects of such arrangements) raise greater concerns about Internet openness than others?

At the heart of the open Internet is the voluntary flexibility of all participants to explore new business and technical arrangements. There was little controversy here until large-volume video providers began to dominate Internet traffic flows. While broadcast, satellite, and cable all provide video distribution services, the Internet is unique in its ability to provide a wide range of interactive, computer-based services. Thus there is some risk that could occur if the Internet is turned into yet another television medium. Video consumes more bits per second, on average, than any other current application, and thus creates a higher load on downstream networks. The best way to deal with this is to let the market decide: if video distributors wish to contribute to the cost of the networks that carry their services to the end users, they should be allowed to do so. If *competitive* networks choose to not expand their video-carrying capacity, in order to make a more economical service available to their customers, they should be allowed to not do so. Internet video distributors should not be allowed to game the regulatory system in order to have, in effect, a subsidized medium with which to compete with other video distributors.

Edge providers are large-scale customers of ISPs

At 55, the Commission asks:

...The court in Verizon also stated that, apart from the service provided to end users, “broadband providers furnish a service to edge providers, thus undoubtedly functioning as edge providers’ ‘carriers.’”¹²⁵ We seek comment on whether this should be identified as a separate service and, if so, how we should define that service and what the regulatory consequences are, if any, of that definition.

Large-scale edge provider facilities are almost always based at data centers where backbone Internet access is provided over high-speed facilities. They also make frequent use of content distribution networks to maintain caches at other backbone locations and at major ISP sites. Edge providers thus subscribe to, or peer with, ISPs, but those are indoor (LAN) connections. Thus the ISP is not really a “broadband provider” in the retail sense, nor is it a carrier *per se*, as its service is only delivered locally, though of course it makes use of carriers (often affiliated) to leave the data center.

This is a classic example of an information service. It thus has little nexus with Title II, and we suggest that the Commission tread very lightly here. Edge providers have access to resources and facilities far in excess of what is available to consumers. There is no evidence of market failure here, even if some providers wish to use the Commission’s regulatory authority to improve their negotiating position in the ISP market. Similarly, at 59, the Commission asks:

Internet Traffic Exchange. The Open Internet Order explained that its rules did not apply beyond “the limits of a broadband provider’s control over the transmission of data to or from its broadband customers.” In other words, the Order applied to a broadband provider’s use of its own network but did not apply the no-blocking or unreasonable discrimination rules to the exchange of traffic between networks, whether peering, paid peering, content delivery network (CDN) connection, or any other form of inter-network transmission of data, as well as provider-owned facilities that are dedicated solely to such interconnection. Thus, the Order noted that the rules were not intended “to affect existing arrangements for network interconnection, including existing paid peering arrangements.” We tentatively conclude that we should maintain this approach, but seek comment on whether we should change our conclusion. Some commenters have suggested that we should expand the scope of the open Internet rules to cover issues related to traffic exchange. We seek comment on these suggestions. For example, how can we ensure that a broadband provider would not be able to evade our open Internet rules by engaging in traffic exchange practices that would be outside the scope of the rules as proposed?

In this case we agree with the tentative conclusion. The peering (traffic exchange) market is functioning well. It is a higher-layer exchange of information between information service providers that may well be outside of the Commission’s jurisdiction, so there is no reason for the Commission to step into that

potential quagmire. Broadband providers will exchange traffic based upon their perceived market requirements. This should be left to a competitive marketplace as well. While the lower-layer broadband telecommunications functions are not sufficiently competitive, and require regulatory intervention as we have noted above, the decisions made by competitive ISPs that operate over these telecommunications services are something else entirely, and there is no one correct answer that the Commission could or should enforce. Rather, the Commission should resist the siren call of edge providers whose business model would depend on regulatory intervention to succeed. This does not exclude the possibility of action by parties other than the FCC if such non-common-carrier activity is conducted in a manner that violates other laws, such as antitrust or restraint of trade. The Commission's nearly-exclusive authority is over interstate common carriage; other agencies and courts have authority over other matters.

Edge providers should not be singled out for special treatment

The Commission asks at 75:

...Some commenters have suggested that current disclosures provide insufficient information for edge providers. We seek comment on how the existing transparency rule is working and how we can enhance its effectiveness with respect to edge providers. Should we view some categories of edge providers, such as start-up companies, as having distinct needs and, if so, what would be the implications for an enhanced transparency rule?

The implication here is that some customers of ISPs, classified as “edge providers”—and within that class, perhaps some called “start-ups”—are more equal than others, and deserve special regulatory treatment. This is not a simple matter, as the application of a regulatory classification regime to Internet users changes its fundamental nature from a flexible one in which market forces prevail to one in which regulatory classifications are worth litigating over. A far better approach would be to simply fix the underlying problem—lack of competitive access to rivalrous transmission facilities—and allow open competition for the information service provider role. Anything else risks the creation of a Rube Goldberg-style classification morass similar to the one that plagues the PSTN.

Likewise we note that at 97 the Commission states:

...Consistent with the court's ruling, we tentatively conclude that the revived no-blocking rule should be interpreted as requiring broadband providers to furnish edge providers with a minimum level of access to their end-user subscribers. We tentatively conclude that our proposed no-blocking rule would allow broadband providers sufficient flexibility to negotiate terms of service individually with edge providers, consistent with the court's view that we must permit providers

to “adapt . . . to individualized circumstances without having to hold themselves out to serve all comers indiscriminately on the same or standardized terms.”...

Again, this threatens to create a class of “edge providers” who are entitled to special treatment. In this case, because many edge providers are not *directly* connected to end users’ ISPs, a minimum level of access to end users implies regulation of peering and transit. This threatens to create a system resembling the switched access rules that are so hard to administer on the PSTN, even though PSTN calls are discrete events that are relatively easy to measure. Applying end-to-end regulation to connectionless IP traffic flows would make switched access look simple by comparison.

To be sure, it is possible to create telecommunications services that have end-to-end guarantees or a minimum level of access. For example, Broadband ISDN, a subject of much discussion in the 1985-1993 time frame, would have created a polyservice telecommunications network with connection options optimized for voice, video, and data flows. But the market did not accept it at the time, nor did the carriers. Perhaps a new polyservice option may arise within the scope of telecommunications services that the Commission can regulate as such. But that is not the Internet. Instead we suggest that a market for polyservice information service offerings, making use of telecommunications services, be allowed to develop freely.

Specialized services are beneficial and fundamentally harmless

The so-called “best effort” services that dominate the public Internet today are not directly suitable, without enhancement, for every application. Some applications that are sensitive to loss, delay, or jitter (collectively, *quality impairment*) perform better when a different type of service is available. These are generally referred to as “specialized services.”

With the notable exception of video, most specialized service applications require relatively little bandwidth. Video is a special case: broadcast-style entertainment, including on-demand entertainment (but not two-way video conferencing), can tolerate several seconds of delay by pre-fetching and buffering the static video stream, and thus does not require the same type of service needed, for example, for good-quality VoIP (which of course cannot be “pre-fetched”). Non-video applications for specialized services are the traditional domain of telecommunications services such as leased lines (special access) and Carrier Ethernet. They are widely used in private networks for many critical applications.

These specialized services should be allowed to continue to thrive and evolve, as this is an area in which technological progress is rapid and largely unpredictable. It is also important to note that this is not a zero-sum game: making specialized services available does not make best-effort services less valuable. They are complementary. Intelligent network design and management can increase the effective

utilization of facilities without impairing service to any one class of user. The Commission thus notes at 60:

We tentatively conclude that we should maintain this approach and continue to closely monitor the development of specialized services to ensure that broadband providers are not using them to bypass the open Internet rules or otherwise undermine a free and open Internet. We seek comment on this tentative conclusion. How can we ensure that the specialized services exception is not used to circumvent our open Internet rules? In addition, should specialized services be addressed within the scope of the “commercially reasonable” rule either as a safe harbor or among the factors for consideration? Should the Commission define “specialized services”?

An implication of this view is that the “best effort” Internet is, policy-wise, the highest priority service, even though it is the one that runs at the lowest actual priority. But this may only be temporary. A polyservice Internet could offer superior performance without significant harm, as it would isolate rate-based streaming flows from those that behave cooperatively (e.g., TCP). The most effective way to protect the public’s interests is via competition, to ensure that different ISPs, whose management practices and service offerings may differ, all have access to essential facilities: the lower-layer transmission media. It would also be unwise to define “specialized services” at this time because doing so would limit the future development of information services, which evolve much more rapidly than telecommunications and should not be subject to the same regulation as telecommunications services. Even specialized telecommunications services should be encouraged to evolve, not *defined* (a word which fundamentally means “limited”).

ISPs should not need to go to great lengths to create meaningless metrics

In the historic PSTN, carriers had to regularly file detailed reports, documenting both their financial status and their service records. The PSTN, once subject to strict Title II regulation, has been largely deregulated, and most carriers are no longer required to file the type of reports that might enable the Commission to detect abuses of monopoly power. Yet in 83, it asks:

In light of these concerns, we tentatively conclude that we should require that broadband providers disclose meaningful information regarding the source, location, timing, speed, packet loss, and duration of network congestion. We seek comment on this tentative conclusion, including on how to implement it in a practical manner that provides meaningful information to end users, edge providers, and other stakeholders without causing undue burden on broadband providers...

These are metrics that might make sense in a predictable PSTN, but make little sense for the best-effort Internet. They would be impractical to measure. Congestion, after all, is an *intrinsic* part of TCP operation. The location of congestion is very hard to identify, as multiple networks peer with each other and IP packets, being connectionless, are discarded silently. The prospect of instrumenting the best-effort connectionless Internet more rigorously than the PSTN itself is technically nonsensical, is probably outside of the Commission's authority, and would impose unwarranted costs on ISPs without providing any meaningful information to anyone. Again, competition is a better approach to ensuring that information that meaningfully distinguishes among the options available to consumers—not just about network congestion—is available.

Blocking and differentiation are both a vital part of Internet management

The Commission states at 89:

...The D.C. Circuit acknowledged the validity of this policy rationale for the no-blocking rule adopted in the Open Internet Order, but vacated the rule because it found that the Commission had failed to provide a legal rationale under which the prohibition would not impermissibly subject broadband providers to common carriage regulation. To address the ongoing concerns with the harmful effects that blocking of Internet traffic would have on Internet openness, we propose to adopt the text of the no-blocking rule that the Commission adopted in 2010, with a clarification that it does not preclude broadband providers from negotiating individualized, differentiated arrangements with similarly situated edge providers (subject to the separate commercial reasonableness rule or its equivalent)....

The question raised in this paragraph is at the heart of the current controversy: can individual arrangements be negotiated? The only possible answer here is the affirmative; individual arrangements not only *can* be negotiated, they *must* be negotiated. The Internet is not common carriage *per se*; it is a set of voluntary agreements among providers to exchange traffic for their mutual benefit, and the mutual benefit often includes the exchange of money. Balancing the value of the traffic (i.e., to the broadband service provider's customers) against the price of bulk service to a commercial customer (i.e., an edge provider), a service provider arrives at some kind of arrangement. To regulate this or require strict tariffs based upon rigid classifications of edge providers vs. other customers would be to impose the worst characteristics of PSTN regulation upon a system that is already working perfectly well and does not need any additional regulation.

We further note that Section 706, as we have explained, lacks the teeth to enforce this type of arrangement. Title II common carriage should apply to basic telecommunications, enabling users

including edge providers, wherever they are located, to have a wide choice of ISPs. But the Internet itself should remain market-based. To the extent that service providers engage in undue discrimination, or favor their affiliates over others, antitrust and other remedies can be employed.

At 94, the Commission adds:

A person engaged in the provision of fixed broadband Internet access service, insofar as such person is so engaged, shall not block lawful content, applications, services, or nonharmful devices, subject to reasonable network management.

The blocking rule itself also should not be instated (we note that it cannot be reinstated as it was never fully in effect before being overturned) at the Information Service layer. Ordinary network management practices involve substantial amounts of blocking. This is not done to deny access to *desired* services, though that is sometimes the result. Rather, it is done to protect the network against abuse, such as spam. We note the existence of the informal “mutually assured destruction” rule. This prohibits ISPs from providing service to spammers. If an ISP allows spam, other ISPs block *all* of its traffic, including that of its non-spammer customers. This makes being a spam-friendly ISP a poor proposition. Furthermore, if an ISP does not block a spammer ISP, then *it* is deemed spam-friendly and it too is blocked. Thus the quarantine applies to all traffic originating from spam-friendly and spam-tolerant ISPs. Because of this, most spam is sent from botnets and compromised hosts, not directly. If blocking were banned, then spammers would be able to dramatically increase the volume of traffic they send. Other security problems could also be worsened.

Some parties will discount this on the grounds that spam blocking should be allowed as reasonable network management. We agree that it should; however, we also note that under the CAN-SPAM Act, spamming in the United States is perfectly legal, provided that the spammer makes certain meaningless gestures that neither compromise the actual anonymity of the source nor provide meaningful relief. Thus a no-blocking rule will likely be first exercised by legal spammers, challenging mutually assured destruction, and lead to further litigation. Another example is the use of Distributed Denial of Service attacks based on amplification via DNS or NTP services. Normally these services are both beneficial and even necessary to the operation of the Internet and its many services, but when intentionally abused, such services must be blocked as a countermeasure. But even that is not the only risk. There are many small ISPs whose service is intentionally filtered, to provide a limited content that a specific, usually religious, community prefers. Would these be allowed under a no-blocking rule? What about streaming services that do not cooperate, and whose streams overwhelm the traffic capacity of ISPs or subscribers? Rural and wireless ISPs in particular have limited capacity, and may have to block low-value high-volume traffic in order to permit other traffic to flow. This should be a commercial decision, not a regulatory one.

At 104, the Commission notes:

...For example, a typical end user may reasonably expect the ability to access streaming video from any provider, place and receive telephone calls using the VoIP service of the end user's choosing, and access any lawful web content. Under this approach, a broadband provider that satisfies these and other reasonable expectations would be in compliance with the no-blocking rule...

We note that the first two of the three applications listed, streaming video and VoIP, are examples of television and telephone services, both of which existed before the Internet, both of which have alternative delivery media subject to specific Commission regulation, and neither of which is well-suited to the best-effort delivery that nominally characterizes the Internet. And this illustrates an issue with the direction of this inquiry in general: is the Internet really going to become just another television medium with incidental voice and other usage, or will it be able to continue to provide the many other applications, which almost always use less capacity than television, that cannot be provided over other media?

Common carriage and absolute bits-are-bits nondiscrimination should be the rule for the underlying telecommunications service that should be unbundled from Internet access. The vertically-integrated model fails here; owning transmission media should not be grounds to discriminate, but information services are not common carriage and should not be treated as such.

Priority is an oversimplification

Much of the argument over "neutrality" uses terminology such as "fast lane" and "priority":

Alternatively, we seek comment on whether we should adopt a no-blocking rule that either itself prohibits broadband providers from entering into priority agreements with edge providers or acts in combination with a separate rule prohibiting such conduct. As discussed below, the record in this proceeding reflects numerous public concerns about the potential for priority agreements to harm an open Internet. How could we address such concerns in the context of the no-blocking rule? If the Commission were to proceed down this alternative path, how should the Commission define "priority"? [at 96]

Similarly:

Some have suggested that the Commission go even beyond the requirements of the Open Internet Order to impose flat bans on pay-for-priority service. We seek comment on these suggestions, including whether all pay-for-priority practices, or some of them, could be treated as per se

violations of the commercially reasonable standard or under any other standard based on any source of legal authority. [at 138]

As we noted above, neither “fast lane” nor “priority” is a useful way to either describe or approach the problem. While simple prioritization is one way to provide enhanced QoS to certain low-volume applications—for instance, it is standard practice for many VoIP services—a true polyservice network will have a more complex tool kit. Again, the best solution is to let the telecommunications service be separated from the information service so that the market and technology can sort out the answer. Polyservice networks of the future may offer priority *up to* a contracted-for level, for instance, while blocking or providing discard-eligible bottom-priority service above that level. Polyservice networks could also set an end-to-end cap on loss, delay, and jitter, and even move some packets to a *lower* priority if that would still meet the contracted-for end to end objective, in order to make room for others to meet their objectives.

A packet-switched network, then, is necessarily a *trading space*. It provides an opportunity for participants to trade local delay, jitter, and loss in order to meet global objectives. And such trading can be complex, which is one reason why today’s public Internet does not yet engage in it except by accident, by making use of indirect methods such as content delivery networks and additional peering points. This does not simply reduce to “priority,” and any attempt to impose regulatory controls based on such a blunt instrument will result in unwanted consequences.

Competition is woefully inadequate

The Commission asks (at 47):

We also seek comment on the state of competition in broadband Internet access service, and its effect on providers’ incentives to limit openness. We seek comment on the appropriate view of whether broadband services with substantially different technical characteristics are competitive substitutes. For example, how should we regard the ability of DSL service with speeds of, for example, 3 Mbps downstream and 768 kbps upstream to constrain conduct by a provider of high-speed broadband with speeds of, for example, 25 Mbps downstream and 3 Mbps upstream (or higher)? How should we regard the geography of broadband competition? From an end user’s point of view, do national practices or market shares have any impact on edge providers, without regard to the definition of a geographic market?

For a large share of the population and many current applications, a 3M/768k service is perfectly adequate. It supports web browsing, email, file transfer, system upgrades, gaming, voice, audio streaming, and many other applications. It also supports streaming video in standard definition, including

web-browser video. It is not a substitute for cable service, and its file transfers will take longer than higher-speed services, but it is still very much an enabling capability. But whether it is viewed as fully competitive in the marketplace is a separate question. It is more of a partial substitute, as the higher-end service will naturally have market advantages.

But the key issue here is not speed but market structure. The plain fact of the matter is that widespread competition at the physical layer is unachievable given the real-world costs of deploying and operating the components of different physical transmission infrastructures (boxes, wires, rights of way, conduit, poles, radios, etc.). These create structural barriers to competition that cannot be relieved by regulation, the market, or changing the laws of physics. The majority of Americans have access to one or two wireline broadband service providers. Wireless service is usually not a substitute at all—not only is it slower, but with CMRS-based data plans typically costing about \$15/GB, and a typical wireline user transferring around 50 GB/month, the numbers simply don't add up. Wireless broadband is a useful tool for mobility, for checking email, for casual web browsing, and for “apps” designed for mobile use, but it does not represent meaningful competition for wireline services. Hence the presence of three or four CMRS data providers does *not* change the wireline monopoly/duopoly market power equation.

There is however no reason for Internet services themselves to be constrained to a monopoly or duopoly, even if the physical layer market has these structural characteristics. Requiring ILECs and cable operators to provide open access telecommunications service to ISPs would allow a competitive ISP market to operate, as it does in many countries. While the Commission's various Orders from 2002-2005 assumed that physical layer competition would arise if open access were cut off, that expectation has been proven wrong by the market. The Commission thus has sufficient empirical evidence, including the very existence of this proceeding and the decision of the Courts in *Verizon*, to prove that the rulings and decisions of previous Commissioners were ill-conceived and should be overturned.

To the extent that ISPs operate on a regional or national scale, their practices might have an impact on edge providers, as well as upon the overall ISP market. Relying on physical-layer competition, on the other hand, ensures inadequate ISP competition. Edge providers' offerings should succeed or fail based upon how the competitive ISP market accepts them, based on true market forces, not regulatory interference.

A simple solution leads to maximum openness

The Commission's Section 706-based approach is tightly wedded to the vertically-integrated model that the Bush-Cheney administration's Commissioners adopted in a massive policy reversal of more than two decades of previous Commission practice. That model has failed. Section 706, by its black letter, is about

encouraging facilities deployment, not managing end-user applications. Title II, on the other hand, provides clear authority to manage the lower layers of the networks on which the Internet can be carried as payload. This leads to a direct answer to the Commission's question at 112:

Alternatively, we also seek comment on whether the Commission should adopt an alternative legal standard to govern broadband providers' practices. How can we ensure that our proposed rule sufficiently protects against harms to the open Internet? How would the rule we propose today change if the Commission were to rely on Title II (or other sources of legal authority) to adopt rules to protect and promote Internet openness? We seek comment on how the goal of the proposed rule—to prevent those broadband provider practices that limit Internet openness—could best be achieved.

The best way to ensure that broadband *facilities* providers do not harm the open Internet is to require them to provide common carriage to any and all ISPs on the same terms that they use for their own Internet service offerings. This rule was applied to telephone carriers in 1980 in the *Computer II* ruling and should be reinstated today, expanded to apply to all providers of rivalrous facilities (i.e., incumbent telephone and cable companies and spectrum-holding wireless carriers). While there is a modicum of competition at the facilities level, the fact is that natural-monopoly effects still apply, especially to wireline; and if there were true competition, there would be no debate today over "neutrality." The Commission is discussing how to regulate monopoly information service providers when it should be discussing how to demonopolize the provision of information services. Thus at 116:

Sound public policy requires that Internet openness be the touchstone of a new legal standard. Accordingly, we tentatively conclude that the Commission should adopt a rule requiring broadband providers to use "commercially reasonable" practices in the provision of broadband Internet access service. Our proposed approach is both more focused and more flexible than the vacated 2010 non-discrimination rule. It would prohibit as commercially unreasonable those broadband providers' practices that, based on the totality of the circumstances, threaten to harm Internet openness and all that it protects. At the same time, it could permit broadband providers to serve customers and carry traffic on an individually negotiated basis, "without having to hold themselves out to serve all comers indiscriminately on the same or standardized terms," so long as such conduct is commercially reasonable.

Again, while flexibility is a good thing, it should be left to ISPs to be as flexible as they want to be, while the physical-facilities providers should offer an open common carrier service, on the same or standardized terms, to enable competition among ISPs. That common carrier service would be content-neutral, though its traffic management options could be negotiated. Carrier Ethernet, for instance, offers a variety of

options that can still be telecommunications service and carry monoservice or polyservice traffic. That reinforces our view about the issue raised in 118:

The core purpose of the legal standard that we wish to adopt, whether the “commercially reasonable” standard or another legal formulation, is to effectively employ the authority that the Verizon court held was within the Commission’s power under section 706. In essence, the court upheld the Commission’s judgment that (1) section 706 grants substantive power to the Commission to take actions...

As we noted above, the Commission’s Section 706 authority, upon further review, is likely to be even narrower than what was approved by the D.C. Circuit in *Verizon*. But even if that were not the case, the Commission should not attempt to judge what is commercially reasonable *within* the Internet itself. The Internet is a complex, fluid, and still-emerging phenomenon. Its performance is not easy to measure; it’s not even easy to decide what “performance” means in anything other than very specific circumstances. Burst speed tests, for instance, are far from accurate; they encourage large buffers (“bufferbloat”), which improve speed test performance while degrading essentially everything else. The Commission can choose to regulate all of the Internet, fundamentally changing its nature and opening up a Pandora’s box of litigation, or it can choose to regulate only the underlying telecommunications network facilities and services, which provide critical access *to* the Internet, as strictly as needed to restore a competitive market and minimize the potential for abuse.

The legal standard, then, should be based on the “simple, specific rule” we propose above: the lower layers of access networks that control rivalrous facilities or spectrum should be made available as regulated common carriage, at just and reasonable rates. Given that the traditional TCP/IP model treats the Internetwork Protocol as the layer above the underlying “network,” this is real network neutrality, *but not internetwork neutrality*. Higher layers and backbone networks need not be regulated. The D.C. Circuit noted in *Verizon* that the Computer II regime inspired the Telecom Act’s definitions. This was a very strong hint that restoring and updating the Computer II framework would be acceptable. Then, given the ability of competitive ISPs to enter markets freely, there would be sufficient competition to ensure that the Internet remains open. We thus oppose any regulation of the Internet itself, but note that if IP (the protocol *per se*) is used to deliver a telecommunications service (such as PacketCable), or is the lowest layer at which an access service is made available to information service providers, then that IP-based service is or is part of a telecommunications service. The simple presence of IP (the protocol *per se*) is not what distinguishes telecommunications services from information services.

That leads to the answer to the question at 121:

...Should the Commission adopt a rule that prohibits unreasonable discrimination and, if so, what legal authority and theories should we rely upon to do so? If the Commission ultimately adopts a Title II approach, how should the Commission define the rule in light of the requirements under sections 201 and 202 of the Act?

Parts of Title II, especially sections 201, 202, and 208, should be applied to telecommunications facilities and services, not to the Internet itself. These services should be offered on a nondiscriminatory basis, at just and reasonable rates, subject to the Commission's complaint procedures. Other activities by ISPs—computing activities, including content-related filtering and selective carriage—are best left to a fully competitive market, like the one that, under the auspices of *Computer II*, produced the open Internet. The Commission should leave in place the bright line distinction between telecommunications and information services that the Telecommunications Act envisioned. This distinction will be needed to allow some future technology to arrive, rather than locking in today's Internet as if it were the last word.

A related question arises at 130:

We believe that consumers of broadband access service should have the ability to exercise meaningful choices. How can we factor consumer choice into our analysis of what is commercially reasonable? Should the Commission look for guidance to section 628 of the Act, which makes it unlawful for cable operators and their affiliated satellite cable programming vendors to engage in “unfair or deceptive acts or practices” with certain purposes and effects?

Cable services are subject to regulation concerning content. They are recognized to have considerable monopoly power, and the industry is dominated by large players who often also have interests in programming. While Section 628 provides guidance for dealing with these circumstances, it is only a second-best approach to addressing the Internet, and a distant second at that. Separation of content and carriage as we have suggested above would substantially reduce the ability of ISPs to abuse any affiliations they have with content provision, as subscribers would regain a choice of ISP over the transmission facilities. Section 628 may, however, provide some additional ancillary authority to apply Computer II-like regulation to cable, not that it appears to be required. The Supreme Court's recent ruling in *Aereo* may also broaden protections for over-the-top video providers.

Collusion is a structural issue that needs addressing

Much of the suspicion of “priority” arrangements has to do with a carrier's affiliated information services receiving special (preferential) treatment that is not available to unaffiliated entities.

Non-exclusive, non-affiliated agreements. AT&T has suggested that the Commission exclude from its review of particular practices any agreement between a broadband provider and an edge

provider if the agreement is not exclusive and if the edge provider is not an affiliate of the broadband provider. AT&T explains that subjecting broadband providers to case-by-case scrutiny in such cases “would unnecessarily impede efficient and pro-consumer arms-length commercial dealings.” [at 141]

This concern is understandable, and again was kept in check by the Computer II/III rules when they were applicable. The arrangements suggested by AT&T resemble common carriage so long as “not exclusive” means that it is open to any similarly situated requesting party, not merely a selected few. Were our suggestion for applying common carriage standards to the lower layer to be implemented, the basic transport facilities would be strictly governed by such a rule; applying it to higher layers would then be unnecessary. In the absence of such an arrangement, polyservice offerings, or other specialized offerings to edge providers or anyone else, should be presumed to be commercially reasonable if they are offered on a general basis to any qualified requesting party.

Section 706 should be construed narrowly

The Commission asks at 145:

We also seek comment on how to construe the specific terms and definitions in Section 706. For example, “advanced telecommunications capability” is defined “without regard to any transmission media or technology, as high-speed, switched, broadband telecommunications capability that enables users to originate and receive high-quality voice, data, graphics, and video telecommunications using any technology.” It is clear that broadband Internet access service is such “advanced telecommunications capability,” but we also seek comment on what other broadband-enabled services may fall within the definition of “advanced telecommunications capability.”...

We have noted that Section 706 most likely does not independently grant broad authority to regulate the Internet; it does however address actual telecommunications facilities to the extent that they can provide advanced services. The services that can make use of that facility certainly should extend beyond the existing Internet. Today’s Internet is, essentially, a prototype internet, the small-i version meaning any set of voluntary agreements among network operators to exchange traffic for their mutual benefit.

Today’s TCP/IP Internet was designed around a monoservice model using protocols designed in the 1970s for a private (Department of Defense) network. If advanced telecommunications capabilities are made available on a common carrier basis, then additional internets, offering specialized services or capabilities that could include, for instance, enhanced security, multimedia support, and mobility, could arise.

Private internets already exist. Interisle manages PSnet, a private network operated on behalf of the Metro Boston Homeland Security Region. It supports the public safety (police, fire, and emergency medical service) needs of nine municipalities centered around Boston. It uses Ethernet and TCP/IP protocols, and standard Internet-compatible equipment, but is isolated from the public Internet. It essentially interconnects the municipal networks in its service area, making it a small-i internet.

PSnet is primarily operated using private microwave links and municipal fiber. A few links are tunneled through cable modems. Some of its client municipalities do make use of leased circuits provided by the local ILEC (Verizon). While it would seem natural for Verizon to make its advanced capabilities available for law enforcement, fire protection, emergency medical service, and other public safety and homeland security applications, Verizon has instead notified them that its existing services will be terminated by 2016. While PSnet serves a geographically-compact region that can be linked via private microwave, and PSnet's bandwidth requirements do not exceed its capabilities, we see no reason why this or other networks, public or private, should not have access to the same transmission facilities that the carrier's own ISP has. Prior to 2005, they generally did. But there is no private offering of FiOS fiber, and raw DSL has been withdrawn. Frame Relay and ATM are grandfathered and are being discontinued, while Carrier Ethernet is only offered as private carriage on a case-by-case basis. Special Access still exists but the carrier appears to be ignoring its Carrier of Last Resort obligations. PSnet is just one example of a private internet; Section 706, if anything, should be used to encourage the development of advanced telecommunications (lower-layer) capabilities that can be used by these networks as well as by the big-I public Internet, as it exists today and as it evolves.

To further illustrate relevant issues, consideration is being given by the PSnet stakeholders to utilize PSnet as a means of connecting municipal networks and individual departments directly to the primary carrier hotels within the region as a means of improving reliability and survivability of access to Internet services while providing greater choice of ISPs to the municipalities. This would leverage unique capabilities inherent in the design and deployment of PSnet to provide capabilities not offered today by any traditional telecommunications (or cable) providers. As municipalities and other users become increasingly dependent on "cloud services," the ability to establish highly resilient access to multiple ISP peering points is an important objective that should be encouraged by the Commission and overall telecommunications policy.

Title II should be applied, but narrowly

As we noted above, the basic telecommunications function of access networks is where Title II common carriage would be appropriate. Thus when the Commission asks at 148:

We seek comment on whether the Commission should rely on its authority under Title II of the Communications Act, including both (1) whether we should revisit the Commission's classification of broadband Internet access service as an information service...

we answer in the clear affirmative. Broadband access *to* the Internet—the lower-layer service that links end users to ISP sites—should be Title II common carriage, a telecommunications service, even though its payload, the Internet itself, should not be. The layered approach taken in *Computer II* was appropriate then and would be appropriate now; the “beads on a string” model of vertical integration has failed. In contrast, that paragraph continues:

...and (2) whether we should separately identify and classify as a telecommunications service a service that “broadband providers . . . furnish to edge providers.”

Edge providers generally locate their servers in data centers where Internet service is competitively available. ISPs exchange traffic with one another; if an edge provider does not like the deal it gets from one ISP, it can go to another. If the edge provider's ISP cannot obtain the peering arrangement that the edge provider wants, then the market is sending a strong signal. In a competitive market, the value of an edge provider's service is weighed against the cost of carrying it. To regulate peering would break the delicate balance that exists and replace it with a clumsy morass of what are effectively switched access charges and litigation; this would benefit almost no one (except perhaps the participants' attorneys). Edge provider activities are really more akin to computing than telecommunications. The Commission should not risk a challenge to its authority by stepping into an area that might be more appropriately supervised by the Federal Trade Commission or the Department of Justice, to the extent that any provider abuses its market power.

Thus our answers to these questions in 149 follow:

We now seek further and updated comment on whether the Commission should revisit its prior classification decisions and apply Title II to broadband Internet access service (or components thereof). How would such a reclassification approach serve our goal to protect and promote Internet openness? What would be the legal bases and theories for particular open Internet rules adopted pursuant to such an approach? Would reclassification and applying Title II for the purpose of protecting and promoting Internet openness impact the Commission's overall policy goals and, if so, how?

The Commission should certainly revisit its prior classification divisions and apply Title II to the telecommunications component of services that provide access to the Internet. As we noted above, we suggest that this service be bounded at the layer below the one that supports the Internet's IP layer, the

one referred to as the “network”;²⁴ for example, DSL (to the ATM, Ethernet, or Frame layer), PON, and DOCSIS would be Title II telecommunications. If the lowest layer offered to competitive providers is IP, then that local use of IP would be telecommunications, to the regional aggregation point where it is handed off. This should allow third-party ISPs to operate their own IP networks encapsulated within the lower layer (even IP) services provided by the owners of rivalrous facilities. But the Internet service itself should not be covered by Title II.

We are very strong advocates of the open Internet, but recognize that it is more a function of competition than of regulation. The layered approach protects Internet openness by allowing open competition among any number of ISPs over the naturally-constrained number of broadband transmission facilities available to reach mass-market subscribers. The legal basis of this approach is clear—it is the model that the framers of the Telecommunications Act had in mind when they defined “telecommunications service” and “information service.” The D.C. Circuit implicitly recognized the validity of this approach in its comments in the *Verizon* case:

Tracking the *Computer II* distinction between basic and enhanced services, the Act defines two categories of entities: telecommunications carriers, which provide the equivalent of basic services, and information-service providers, which provide the equivalent of enhanced services.²⁵

...

Pursuant to the Act, and paralleling its prior practice under the *Computer II* regime, the Commission then classified Digital Subscriber Line (DSL) services—broadband Internet service furnished over telephone lines—as “telecommunications services.” See *In re Deployment of Wireline Services Offering Advanced Telecommunications Capability*, 13 F.C.C.R. 24012, 24014, 24029–30 ¶¶ 3, 35–36 (1998) (“*Advanced Services Order*”). DSL services, the Commission concluded, involved pure transmission technologies, and so were subject to Title II regulation. *Id.* at 24030–31 ¶ 35. A DSL provider could exempt its Internet access services, but not its transmission facilities themselves, from Title II common carrier restrictions only by operating them through a separate affiliate (i.e., a quasi-independent ISP). *Id.* at 24018 ¶ 13.²⁶

Reclassification of the network, then, supports the Commission’s stated goal of openness. It restores the competitive nature of the Internet access business, separates content from carriage, and allows any provider access to essential telecommunications facilities on uniform terms. Limiting reclassification to the lower layers affirms the goal by allowing unfettered evolution of information services, not locking in a snapshot of the technology of the Internet as it appeared to exist at any one point in time. Because the common carrier circuits will not involve peering or intercarrier connections, they will be easy to

²⁴ This use of “network” dates back to the earliest descriptions of IP; for example, IEN54, from 1978, refers to the “local network protocol” in its Fig. 1 and elsewhere. That use of “local” does not correspond to the term “local area network” but instead refers to the layer below IP. We are suggesting that the typical geographic boundary of this “network” be on a metropolitan-area scale.

²⁵ *Verizon* at 8.

²⁶ *id.* at 9.

administer and monitor for compliance; as common carriage, essentially any and all interference with the payload (content beyond the protocol header) can be treated as a *prima facie* violation. Activities such as deep packet inspection may or may not have a place in some ISPs' networks but would be strictly forbidden within the common carrier access service itself.

Reclassification is fully justified by the factual record

We note that the phrase “network neutrality” was not even coined until after the Commission, in 2005, granted general forbearance from its Computer II regulations. There was simply no problem to contend with when common carriage was generally available. The Commission's experiment since then with vertically-integrated deregulation has been a failure; complex workarounds within the higher layers of networks are neither appropriate nor necessary. The Commission asks at 150:

What factors should the Commission keep in mind as it considers whether to revisit its prior decisions? Have there been changes to the broadband marketplace that should lead us to reconsider our prior classification decisions?

Clearly there have been. Both the *Cable Broadband Order* and the *Wireline Broadband Order* suggested that facilities owners would *voluntarily* provide access to third-party ISPs. The technical problems associated with third-party ISP access are more complex in cable, but the *Cable Broadband Order* noted²⁷ (with more detail in its footnote 65 than in the text):

Multiple-ISP access is occurring in the marketplace and in trials however, using various routing techniques.⁶⁵

⁶⁵ See AOL Time Warner Jan. 22, 2002 Ex Parte at 5; Excite@Home Aug. 17, 2001 Ex Parte, Attachment at 20. In addition to source-based and destination-based routing, other possible routing techniques include Point to Point Protocol over Ethernet (“PPPoE”) and Layer Two Tunneling Protocol (“L2TP”) tunneling. PPPoE and L2TP are tunneling protocols that enable a Point to Point Protocol (“PPP”) session between the subscriber and the specified ISP. A tunnel is a virtual dedicated connection between two points in a network. Tunneling allows data to traverse through an “intervening” network of a different protocol and works by encapsulating data from one protocol format into another protocol format. PPPoE enables PPP to run over bridged networks, and L2TP enables PPP to run over routed networks. See Letter from Emy Tseng, MIT, *et al.* to Magalie Roman Salas, Secretary, FCC in CS Docket No. 00-30 (May 1, 2000), Attachment at 16-17.

While cable systems were not designed for multiple-ISP access the way DSL was (since it was *required* for DSL until 2005 and never required of cable), the market position of cable modems has changed. ILECs are abandoning their wireline networks in some areas. They have essentially moved from being competitors with cable to partners, allowing cable to have a monopoly on wireline access to customers in areas where the ILEC does not have fiber-based facilities and does not choose to build them. Verizon, for

²⁷ *Cable Broadband Order* at 15.

instance, has stated openly that it does not intend to expand its FiOS footprint beyond areas already designated, but will instead partner with cable, and only provide wireless service in such areas. The cable industry, in turn, has generally abandoned its own wireless plans, selling most of its licenses to Verizon, and is helping market the carriers' wireless services. For instance, Comcast sent out a mailing in June 2014 that offers its cable subscribers a \$100 rebate on the purchase of a new smartphone from Verizon Wireless. There is clearly less competition between the two industries than was anticipated at the time of the *Cable Broadband Order*.

The *Wireline Broadband Order* was even more clear, and (noting additional quotes on point referenced above) thus even less prescient:

Given the nature and history of the broadband Internet access services industry, we expect that wireline broadband transmission will remain available to ISPs and others without any *Computer Inquiry* requirements. Incumbent LECs have represented that they not only intend to make broadband Internet access transmission offerings available to unaffiliated ISPs in a manner that meets ISPs' needs, but that they have business incentives to do so. For example, Qwest offers a tariffed wireline broadband DSL service that enables hundreds of independent ISPs to serve end-user customers over Qwest's broadband facilities. Regardless of the outcome of this proceeding, Qwest has stated it will continue to make available a DSL offering that will enable consumers to reach unaffiliated ISPs because consumers demand the choice, and meeting that demand makes its product more attractive. SBC previously entered into a memorandum of understanding with a trade association representing nearly 300 members of the Internet industry, including many independent ISPs, committing to negotiate private commercial arrangements with unaffiliated ISPs for broadband Internet access.²⁸

Either the carriers misrepresented their intentions, or these predictions were disingenuous or simply wrong. ILECs did in general continue to offer DSL on a detariffed basis for a short period of time, typically via three-year contracts; in some cases these were renewed for an additional three years. But in most cases these agreements were only for older, lower-speed DSL products, not the higher-speed DSL services used by the ILECs' own ISPs, or for fiber-based successor services. The aforementioned agreement with a trade organization did continue for several years, but the most recent offering removes (lower layer raw) DSL as a service, allowing independent ISPs to resell only AT&T's own Internet service, but not allowing them to use AT&T's facilities for access to their own.²⁹ Thus the benefits of openness and competitive service access are no longer available; ISPs that participate are essentially

²⁸ *Wireline Broadband Order* at 74.

²⁹ This newer offering is AT&T "Business DSL", described at http://serviceguidenew.att.com/sg_CustomPreviewer?attachmentId=00PC000000RgzOiMAJ

reduced to the level of commissioned sales agents for AT&T's captive ISP. Verizon has withdrawn DSL from some areas, even its own DSL, and is not generally offering it to independent ISPs. Qwest did for some time offer competitive ISPs nominal access to its services, but numerous abuses led to their general abandonment by the ISP industry.³⁰ Because DSL detariffing was permissive, some rate-of-return carriers have retained their DSL tariffs, primarily under NECA Tariff 5, for reasons apparently related to universal service funding.

Thus the marketplace has not developed as envisioned, and the Commission is fully justified in reconsidering its various Orders of 2001-2005, which collectively shut off independent ISPs' wholesale access to transmission facilities. Nor are there serious technical barriers to moving back to a layered approach. The Commission continues in 150:

To what extent is any telecommunications component of that service integrated with applications and other offerings, such that they are "inextricably intertwined" with the underlying connectivity service?

The lower layers are not themselves inextricably intertwined with applications. However, the *IP* layer sometimes *is*. The TCP/IP stack was designed to be used *atop* a "network," such as a packet-switched service or leased facility, hence it is an "internetwork" from IP on up. Separating IP (except when used for local purposes, when it may transparently encapsulate another instance of IP, or as the multiplexing function of a non-Internet service, such as Voice Using IP) from the application could lead to serious technical issues. We reiterate for example how the Domain Name System operates within the application layer, and how packet redirection is performed using application labels (URLs). Thus we strongly suggest that the boundary of the telecommunications component be drawn below the Internet Service Provider's IP layer.

We do not however call for *every* ISP to have its lower-layer facilities or services unbundled. ISPs who manage to lease transmission facilities do not have the same market power as the facilities owners. Wireless ISPs who operate using unlicensed or non-exclusively-licensed frequencies almost always use systems that do not easily separate out the IP layer, and they too do not use rivalrous facilities that provide market power. Hence their operation should not be affected at all by these rules.

The open Internet is today's free press

At 159, the Commission raises important constitutional issues:

³⁰ One Interisle principal served as expert witness for the Attorney General of New Mexico in an investigation of Qwest's practices, which found a majority of the state's independent ISPs unhappy with the DSL service delivered in the late 2000s; most thus stopped using it. We are not aware of whether or not CenturyLink, as Qwest's successor-in-interest, even continues to make the service available at all.

Finally we seek comment on other legal limitations and barriers to adoption of the rules we propose today, including First Amendment and Due Process considerations. In the Open Internet Order, the Commission concluded that “broadband providers typically are best described not as ‘speakers,’ but rather as conduits for speech,” and that the open Internet rules therefore did not implicate broadband providers’ First Amendment rights. The Commission also found that even if the rules “did implicate expressive activity, they would not violate the First Amendment” because they would advance an important government interest—“ensur[ing] the public’s access to a multiplicity of information sources and maximiz[ing] the Internet’s potential to further the public interest”—without burdening “substantially more speech than is necessary.” We seek comment on these findings.

The problem with the Commission’s previous conclusion is that it attempted to find a common answer to two different questions. The rights and responsibilities of a carrier are not the same as the rights and responsibilities of a carrier’s customer (speaker) or, in a newer context, a content owner. In the context of the telephone network, the distinction is obvious. Telephone companies are traditional common carriers, and they have no responsibility for, nor the right to control, what is said on their lines by speakers.

The Internet is different because it adds a third role. One role, the telecommunications function, can and should be viewed as common carriage (at least in the access network), although at this time it is not. Another, that of the end user or edge provider, is clearly content, and should be treated as speech or as *the press*, constitutionally protected but subject to the laws of copyright, libel, etc.

The third role, that of the information service provider, is more problematic. It is not common carriage, but neither is it end user content. The key statutory guidance here seems to come from 47 U.S.C. 230(c) (1):

TREATMENT OF PUBLISHER OR SPEAKER.—No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.

This law predates public access to the Internet. It was passed at a time when online computer services such as Compuserve and Delphi were more widely used. These were basically computer time-sharing systems with shared applications and data. Section 230(c) (1) gave them protection when they allowed their customers to post information. This created an intermediate stage between full common carriage and full responsibility for content, although online publishers remained responsible for their own content.

By the time the Telecommunications Act had passed, online services had begun to offer Internet access, but were still not pure ISPs. In 1995, large services including AOL, Compuserve, and Prodigy were providing their users with access to Internet email, and were beginning to add web browsing, but they

were still time-sharing services at heart. AOL, the largest online service of the dial-up era, used a proprietary protocol (“P2”) to display text and graphics on a customer’s computer; that computer had to use the client program, which was typically distributed on a floppy disk. The only Internet applications its users could access were those that AOL provided. Later, some time after the Telecommunications Act was passed, AOL began to offer direct TCP/IP access to its subscribers, and it became what we would today think of as a dial-up ISP. (Even later, as that business died off, it became primarily a content provider.) Other online service providers followed a similar trajectory. Eventually the online services were forgotten and the mass market is dominated today by a relatively homogenous model in which ISPs offer an IP-based network monoservice that gives the illusion of unrestricted, neutral access to the global Internet. We suggest that regulations that would find a service such as 1996’s AOL in violation are obviously not what was intended by a law written during a time when such “training wheels for the Internet” services were still growing by leaps and bounds.

Information services are not the same as edge providers, who are essentially speakers. Nor are they carriers. The public Internet evolved because the Computer Inquiries and the Telecommunications Act drew a boundary above the basic carriage function and did not attempt to regulate, or distinguish between, speakers and information services. Indeed large edge providers operate their own networks in order to distribute their content, making them in effect their own ISPs. Just who is and isn’t an ISP, however, is undefined: carriers asking for forbearance from the Computer Inquiries in the early 2000s at times characterized all web sites as “ISPs,” and Verizon claimed that there were thus millions of ISPs, so cutting off access to the ISPs who provided competitive access to mass-market subscribers had only a negligible impact on the total count of what it considered to be ISPs. This was of course a ridiculous semantic trick. But like many subterfuges, it built upon a grain of truth: that there is not a bright line distinction between ISP and content or edge provider. Fluidity and flexibility, even in business structure, have been the hallmark of the open Internet.

Content-based regulation of information service providers thus does raise constitutional issues, as they are users of the carrier function, not carriers themselves. They may be vertically integrated today, but the content and carriage functions are easy to distinguish. There is no barrier to requiring common carriers to provide absolutely content-neutral transport; that is their business. But there is no need to force information service providers to pretend to be common carriers, nor does the Commission have that within its clear authority. It is better to view ISPs and edge providers together as being the electronic version of newspapers, while the local facility owners are the equivalent of newsprint or the postal service (which was, after all, created by the Founding Fathers as a content-neutral medium of newspaper delivery). Adopting a strict Title II common carriage regime limited to lower-layer transport networks

avoids this constitutional controversy while providing the tools necessary to preserve and continue to evolve the open Internet.

Respectfully submitted by Interisle Consulting Group, LLC and its principals

Fred R. Goldstein

A. Lyman Chapin

Charles Wade

Colin Strutt